

Relative fluid stretching and rotation for sparse trajectory observations

Nikolas O. Aksamit^{1,†}, Alex P. Encinas-Bartos², George Haller² and David E. Rival³

¹Department of Mathematics and Statistics, UiT – The Arctic University of Norway, Forskningsparken 1, Tromsø 9037, Norway

²Institute for Mechanical Systems, Swiss Federal Institute of Technology (ETH) – Zürich, Leonhardstrasse 21, Zürich 8092, Switzerland

³Institute of Fluid Mechanics, Technische Universität Braunschweig, Hermann-Blenk-Strasse 37, Braunschweig 38108, Germany

(Received 9 October 2023; revised 10 June 2024; accepted 28 August 2024)

As most mathematically justifiable Lagrangian coherent structure detection methods rely on spatial derivatives, their applicability to sparse trajectory data has been limited. For experimental fluid dynamicists and natural scientists working with Lagrangian trajectory data via passive tracers in unsteady flows (e.g. Lagrangian particle tracking or ocean buoys), obtaining material measures of fluid rotation or stretching is an active topic of research. To facilitate frame-indifferent investigations in unsteady and sparsely sampled flows, we present a novel approach to quantify fluid stretching and rotation via relative Lagrangian velocities. This technique provides a formal objective extension of quasi-objective metrics to unsteady flows by accounting for mean flow behaviour. For extremely sparse experimental data, fluid structures may be significantly undersampled and the mean flow behaviour becomes difficult to quantify. We provide a means to maintain the accuracy of our novel sparse flow diagnostics in extremely sparse sampling scenarios, such as ocean buoy data and Lagrangian particle tracking. We use data from multiple numerical and experimental flows to show that our methods can identify structures beyond existing limits of sparse, frame-indifferent diagnostics and exhibit improved interpretability over common frame-dependent diagnostics.

Key words: atmospheric flows, ocean processes, chaos

† Email address for correspondence: nikolas.aksamit@uit.no

1. Introduction

Experimental methods for spatially and temporally resolving fluid velocities in both large and natural flow domains have improved significantly in recent years. For a wide range of turbulent flows, ground truth measurements of fluid velocity can be measured using high-speed imaging of advected particles and any number of particle image velocimetry (PIV) and Lagrangian particle tracking (LPT) algorithms. For the largest spatial scales and for the sparsest trajectory data, only LPT is suitable and provides a Lagrangian framework with which to extract transport features in the flow. At oceanographic and atmospheric circulation spatial scales, GPS tracking of buoys and balloons typically replaces PIV and LPT imaging approaches. To date, Lagrangian data have provided great insights to our studies of sea ice and ocean dynamics, as well as for meteorologists studying atmospheric behaviours with drifting weather-balloon measurements (Businger, Johnson & Talbot 2006; Leppäranta 2011; van Sebille *et al.* 2018). Coherent structure identification from sparse data, however, has typically relied on a grab-bag of techniques, often tailored to each individual flow, with no unifying metrics that work in all domains.

Historically, lab-based measurement techniques have evolved hand-in-hand with systematic technological advances such as integrated circuits, the laser and most recently the CMOS chip. Starting out from intrusive, probe-based extraction of Eulerian data (fixed-point statistics) through to time-averaged planar field measurements, e.g. PIV, and then most recently to dense time-resolved particle tracking in three dimensions, the availability of said tools has influenced the choice of metrics used to describe the flow in question. As a case in point, the use of Reynolds stresses to describe shear flows has dominated the community since the days of hot-wire anemometry even though such stresses are only a proxy to the coherent structures driving the turbulent processes on hand.

Commonly used approaches for identifying structures in experimental flows are typically frame-dependent. That is, the extracted features will depend on the choice of reference frame of the experimentalist and, thus, violate a fundamental requirement from continuum mechanics for describing *material* fluid behaviour. Material behaviour can be thought of as the features in a flow revealed in a tracer visualisation experiment (e.g. dye or smoke). Although our physical intuition around Lagrangian velocities may be strong, e.g. one can easily imagine a leaf floating downstream on a river surface, extracting physically meaningful diagnostics that describe the material deformation of the surrounding fluid is much more difficult. Indeed, many common and intuitive trajectory metrics are frame-dependent, such as the Lagrangian velocity, looping (Lumpkin 2016), curvature (Bristow *et al.* 2023), complexity measures (Rypina *et al.* 2011) and network-based approaches (Iacobello & Rival 2023), as well as many diagnostics from gridded velocity data such as vorticity (Bernard & Thomas 1993), swirling strength (Zhou *et al.* 1999) and the lambda-2 criterion (Jeong & Hussain 1995).

Although choosing a single common reference frame may appear logical to study flows we understand well, this approach quickly loses its foundation when encountering flows where no *a priori* knowledge of the relevant structures or a natural reference frame are available. Furthermore, truly material behaviour does not change under rigid-body frame changes, regardless of whether a user-preferred frame exists or not (Truesdell & Noll 2004). Thus, Euclidean frame-indifference is a fundamental litmus test for structure identification schemes to actually identify material features even if one always conducts their research in the same reference frame.

Methods that can effectively reveal fluid structures at both high and low trajectory densities, in a frame-independent manner, and for the widest variety of flows would provide a wealth of information for both understanding fluids in observational

environments and as a common ground for comparison with numerical simulations. In order to test sparse diagnostics in very different flows, we quantify the sparsity of a data set by normalising the number of trajectories by the square or cube of a characteristic length scale of the flow ℓ , for two-dimensional (2-D) and three-dimensional (3-D) examples, respectively.

Frame-indifferent (objective) diagnostics that identify Lagrangian coherent structures (LCS) have been developed extensively over the last two decades (see Haller 2023) but their implementation relies on spatial derivatives that are difficult to accurately compute from sparse or unstructured data. To account for this, adaptations have been developed to allow for data sparsity (Lekien & Ross 2010; Rypina *et al.* 2021; Mowlavi *et al.* 2022) as have Green's theorem-based approximations of rate-of-strain metrics from trajectory arrays (Kwok *et al.* 1990). Most notably, Mowlavi *et al.* (2022) compared multiple sparse methods for identifying hyperbolic (stretching) and elliptic (rotating) LCS in the Bickley jet and ABC flow. When initialising particles on a structured grid, Mowlavi *et al.* (2022) were able to accurately cluster particles into elliptic LCS at particle concentrations of $85\ell^{-2}$ and $504\ell^{-3}$ for the Bickley jet and ABC flow ($\ell = \pi$ for both), respectively. They were also able to identify hyperbolic structures from randomly initialised particles at concentrations of $471\ell^{-2}$, and $3870\ell^{-3}$ for the Bickley jet and ABC flow, respectively.

Modern clustering methods provide a complementary frame-indifferent approach to identify regions of fluid with similar fluid particle trajectories (see, e.g. Froyland & Padberg-Gehle 2015; Hadjighasem *et al.* 2016; Schlueter-Kuck & Dabiri 2017). Graph-theory-based clustering algorithms quantify similarity between trajectories themselves and, thus, do not require measurements to be spatially proximal, as is required for gradient-reliant approaches. As such, trajectories can be generated from gridded flow data or observed experimentally, such as by LPT. In contrast to hyperbolic and elliptic LCS, the coherent structures identified with trajectory clustering algorithms have no inherent physical meaning besides a certain trajectory similarity. That is, one cannot use clustering alone to interpret local stretching or rotation rates without *a priori* knowledge of the flow behaviour.

The squared relative dispersion (d^2) is an outlier diagnostic as it is suitable for sparse and randomly oriented trajectories, is also objective (Haller & Yuan 2000), and has been developed with a strong physical foundation. This fluid stretching metric has been used for a number of years, particularly to understand dispersion and mixing by oceanographers in a statistical manner (LaCasce 2008). The ability of d^2 to identify coherent structures, however, is limited as one is forced to initially choose particle pairs, whose relative motion inevitably becomes uncorrelated at an *a priori* unknown temporal horizon (Haller, Aksamit & Bartos 2021).

Current approaches for physically meaningful sparse trajectory diagnostics still rely on a relatively dense field of particles (for a comparison, see, e.g. Mowlavi *et al.* 2022), or well-behaved trajectory array geometries (Lindsay & Stern 2003). In light of the unstructured nature of trajectory data, it is more common for experimental LPT data to be spatially and temporally averaged or assimilated into gridded products prior to structure analysis (Schröder & Schanz 2023). When averaging or projecting time-dependent fluid motions to a single instance in time, such approaches potentially oversimplify transient (non-stationary) flow behaviour that may exist in the underlying Lagrangian data.

Complementing these sparse approaches is the recent development of quasi-objective coherent structure diagnostics, the trajectory rotation average (TRA) and trajectory stretching exponent (TSE) (Haller *et al.* 2021). TSE and TRA calculate stretching and rotation, respectively, for individual particle trajectories with no requirement of

nearby velocity or trajectory data. The authors mathematically proved that under suitable conditions (slowly varying, relatively small mean vorticity) TRA and TSE approximate objective measures of rotation and stretching. Given these flow conditions, TRA and TSE have proven to be advantageous in several extremely sparse geophysical buoy experiments, accurately identifying algae trapping eddies in the ocean (Encinas-Bartos, Aksamit & Haller 2022), as well as predicting Arctic sea ice stretching and breakup events that were missed by other approaches (Aksamit *et al.* 2023).

In the present research, we introduce relative stretching and rotation metrics for individual trajectories that incorporate knowledge of the average translation and rotation of the sampled fluid. By utilising bulk behaviour of concurrent trajectories in a given experiment, we can obtain objective diagnostics of stretching and rotation in unsteady flows with much of the same flexibility of a true single-trajectory method. Our measures of relative stretching and rotation can be seen as a natural synthesis of the objective Eulerian deformation velocity of Kaszás, Pedergnana & Haller (2023) and quasi-objective diagnostics of Haller *et al.* (2021).

In the following, we develop the theory of relative stretching and rotation and provide several examples of performance. The robustness of relative Lagrangian stretching and rotation is displayed by first comparing full resolution structure topology against ground truth LCS boundaries in 2-D and 3-D flows, and then testing the accuracy of the trajectory diagnostics in progressively downsampled data sets. We also display the enhanced ability of relative rotation to distinguish between experimental turbulent flows for extremely sparsely sampled data, when compared with traditional metrics. This superior performance is shown for numerical simulation data and in real-world observations of ocean buoys from the global drifter database and a large-scale LPT wind tunnel experiment.

2. Methods

2.1. Background

Consider a fluid flow with time-varying velocity field $\mathbf{v}(\mathbf{x}, t)$. Infinitesimal fluid particle trajectories can be generated as solutions to the differential equation $\dot{\mathbf{x}}(t) = \mathbf{v}(\mathbf{x}(t), t)$ from some initial position $\mathbf{x}(t_0) = \mathbf{x}_0$. The flow map \mathbf{F} maps fluid particles from their time t_0 -positions to their position at a time t , along the trajectory $\mathbf{x}(t; t_0, \mathbf{x}_0)$,

$$\mathbf{F}_{t_0}^t : \mathbf{x}_0 \mapsto \mathbf{x}(t; t_0, \mathbf{x}_0). \quad (2.1)$$

The measurement of a scalar quantity associated to a fluid particle $P(\mathbf{x}(t))$, such as its temperature, is objective (frame-indifferent) under Euclidean transformations of the form

$$\mathbf{x} = \mathbf{Q}(t)\mathbf{y} + \mathbf{b}(t), \quad (2.2)$$

where $\mathbf{Q}(t)$ is a rigid-body rotation matrix (an element of the matrix group $\text{SO}(n)$ for $n = 2$ or 3) and $\mathbf{b}(t)$ is a time-varying translation vector. That is, in the two reference frames, the scalar quantity remains the same: $\tilde{P}(\mathbf{y}) = P(\mathbf{x})$, where P and \tilde{P} are measured in the original and translated frames, respectively.

Similarly, one can define an objective Lagrangian vector $\boldsymbol{\xi}$ as one that transforms under (2.2) as

$$\tilde{\boldsymbol{\xi}}(\mathbf{y}(t, \mathbf{y}_0)) = \mathbf{Q}^T(t)\boldsymbol{\xi}(\mathbf{x}(t, \mathbf{x}_0)). \quad (2.3)$$

To reveal the ground-truth material fluid structures undergoing significant stretching and rotation, we will rely on the frame-indifferent finite-time Lyapunov exponent (FTLE)

(Haller 2015) and the Lagrangian-averaged vorticity deviation (LAVD) (Haller *et al.* 2016), respectively:

$$\text{FTLE}_{t_0}^t(\mathbf{x}_0) = \frac{1}{2|t-t_0|} \log \lambda_{\max}(\mathbf{C}_{t_0}^t(\mathbf{x}_0)), \quad (2.4)$$

$$\text{LAVD}_{t_0}^t(\mathbf{x}_0) = \frac{1}{|t-t_0|} \int_{t_0}^t |\boldsymbol{\omega}(\mathbf{F}_{t_0}^s(\mathbf{x}_0), s) - \bar{\boldsymbol{\omega}}(s)| \, ds, \quad (2.5)$$

where $\lambda_{\max} > 0$ is the largest eigenvalue of the positive-definite tensor $\mathbf{C}_{t_0}^t = [\nabla \mathbf{F}_{t_0}^t]^T \nabla \mathbf{F}_{t_0}^t$ and $\bar{\boldsymbol{\omega}}$ is the time-varying but spatially averaged vorticity. Although LAVD and FLTE have seen widespread use in studies of atmospheric, oceanic and experimental flows, their direct application to sparse data sets has been hindered by their strong reliance on velocity and flow map gradients.

In the sparse sampling context, where we have no control over trajectory densities or orientations, it would be advantageous to develop a LCS algorithm based on individual trajectories. The primary issue surrounding the development of an objective single-trajectory diagnostic is that one can always pass to the frame of the particle, and in that reference frame, the particle is not moving. As a way around this frame-dependency, Haller *et al.* (2021) recently developed the quasi-objective coherent structure diagnostics

$$\overline{\text{TSE}}_{t_0}^{t_N}(\mathbf{x}_0) = \frac{1}{t_N - t_0} \sum_{j=0}^{N-1} \left| \log \frac{|\dot{\mathbf{x}}(t_{j+1})|}{|\dot{\mathbf{x}}(t_j)|} \right|, \quad (2.6)$$

$$\overline{\text{TRA}}_{t_0}^{t_N}(\mathbf{x}_0) = \frac{1}{t_N - t_0} \sum_{j=0}^{N-1} \cos^{-1} \frac{\langle \dot{\mathbf{x}}(t_j), \dot{\mathbf{x}}(t_{j+1}) \rangle}{|\dot{\mathbf{x}}(t_j)| |\dot{\mathbf{x}}(t_{j+1})|}, \quad (2.7)$$

which provide close approximations to objective stretching and rotation measures *if* the Lagrangian trajectories are analysed in slowly varying reference frames with relatively small mean vorticity.

In the following, we derive the experiment-relative Lagrangian velocity of a fluid particle and show that it is an objective vector. With these relative velocities, we can objectively define the relative stretching and rotation of fluid from sparsely sampled experimental data, with no *a priori* knowledge of the structures being investigated.

2.2. Relative Lagrangian velocity

Suppose we have a collection of Lagrangian particle trajectory observations $\{\mathbf{x}_i(t)\}$, with corresponding velocities along their trajectories $\dot{\mathbf{x}}_i(t) = \mathbf{v}_i(t)$. By linearity of the time derivative, the velocity of the average position of these particles at time t is equivalent to the average of their Lagrangian velocities,

$$\bar{\mathbf{x}}(t) = \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i(t), \quad (2.8)$$

$$\mathbf{v}(t) = \dot{\mathbf{x}}(t), \quad \bar{\mathbf{v}}(t) = \frac{1}{n} \sum_{i=1}^n \mathbf{v}_i(t) = \dot{\bar{\mathbf{x}}}(t). \quad (2.9)$$

Furthermore, define the time-varying moment-of-inertia tensor

$$\boldsymbol{\Theta}(t) = \overline{|\mathbf{x}_i - \bar{\mathbf{x}}|^2 \mathbf{I} - (\mathbf{x}_i - \bar{\mathbf{x}}) \otimes (\mathbf{x}_i - \bar{\mathbf{x}})}, \quad (2.10)$$

where averaging is over the trajectory index i . From our collection of trajectories, mimicking the Eulerian calculations of Kaszás *et al.* (2023), we relate the experiment-based vorticity to the approximate rigid-body rotation of the flow as

$$\bar{\omega}_{rb}(t) = \Theta^{-1} \overline{(\mathbf{x}_i - \bar{\mathbf{x}}) \times (\mathbf{v}_i - \bar{\mathbf{v}})}. \quad (2.11)$$

To avoid cumbersome notation, we drop the subscript i from hereon when it is clear we are referring to a specific trajectory $\mathbf{x}(t; t_0, \mathbf{x}_0)$. Having now calculated the average rotation and translation of the sampled fluid in our experiment, we define the relative velocity along a Lagrangian trajectory as

$$\mathbf{v}_d(\mathbf{x}(t)) = \mathbf{v}(\mathbf{x}(t)) - \bar{\mathbf{v}}(t) - \bar{\omega}_{rb}(t) \times (\mathbf{x} - \bar{\mathbf{x}}). \quad (2.12)$$

The objectivity of \mathbf{v}_d , and scalar diagnostics generated from it, follows from the objectivity of \mathbf{v}_d originally derived by Kaszás *et al.* (2023). That is, by removing the mean rate of translation and rotation of our trajectory observations, we can study objective properties of dynamic fluid structures in our sampled domain.

This is a Lagrangian-observation-based implementation of the Eulerian deformation velocity \mathbf{v}_d originally derived by Kaszás *et al.* (2023). This is, however, distinct from calculating Lagrangian trajectories in time-resolved Eulerian \mathbf{v}_d fields. Trajectories calculated in the Eulerian \mathbf{v}_d do not necessarily trace material flow features, but observed experimental tracers will (Kaszás *et al.* 2023). Instead, \mathbf{v}_d can be used to objectivise frame-dependent Eulerian diagnostics in each snapshot. Here, by estimating rigid-body translation and rotation from passive Lagrangian measurements, the relative Lagrangian velocity \mathbf{v}_d provides a means to objectivise trajectory diagnostics relative to the sparse data available without *a priori* knowledge of the underlying flow.

2.3. Stretching

For steady flows, Haller *et al.* (2021) showed that cumulative fluid stretching and compression normal to trajectories can be quantified using $\overline{\text{TSE}}$ (2.6). In slowly varying flows, this trajectory stretching and compression approximates the true material deformation with the difference being a function of $\partial_t \mathbf{v}(\mathbf{x}(t), t)$. Furthermore, TSEs were shown to faithfully highlight hyperbolic LCS and reproduce the dominant features of FTLE fields.

Here, we seek a measure of fluid stretching that is frame-indifferent in unsteady flows and has minimal dependence on trajectory concentration, thus making it suitable for extremely sparse sampling in many natural environments. To do this, we introduce the stretching of fluid parcels relative to the bulk behaviour of our flow through the stretching of relative Lagrangian velocity vectors. Formally, the *relative* trajectory stretching exponents (rTSE and $\overline{\text{rTSE}}$) from time t_0 to time t_N can be written as

$$\text{rTSE}_{t_0}^{t_N}(\mathbf{x}_i) = \frac{1}{t_N - t_0} \log \frac{|\mathbf{v}_d(\mathbf{x}_i(t_N))|}{|\mathbf{v}_d(\mathbf{x}_i(t_0))|}, \quad (2.13)$$

$$\overline{\text{rTSE}}_{t_0}^{t_N}(\mathbf{x}_i) = \frac{1}{t_N - t_0} \sum_{j=0}^{N-1} \left| \log \frac{|\mathbf{v}_d(\mathbf{x}_i(t_{j+1}))|}{|\mathbf{v}_d(\mathbf{x}_i(t_j))|} \right|. \quad (2.14)$$

The rTSE diagnostic provides the stretching exponent for the relative Lagrangian velocity vector from time t_0 to t_N , whereas $\overline{\text{rTSE}}$ is a cumulative measure of all stretching and contraction that occurs in the same time window. Given the objectivity of the relative

Lagrangian velocity \mathbf{v}_d defined in (2.12), then rTSE and $\overline{\text{rTSE}}$ are also objective. That is, because $\hat{\mathbf{v}}_d(\mathbf{y}(t)) = \mathbf{Q}^T \mathbf{v}_d(\mathbf{x}(t))$, we have $|\hat{\mathbf{v}}_d(\mathbf{y}(t))| = |\mathbf{v}_d(\mathbf{x}(t))|$ and rTSEs do not change between reference frames. This provides a formal and fully objective extension of the quasi-objective TSEs from Haller *et al.* (2021) that works in unsteady (time-varying) flows.

2.4. Rotation

For steady flows with negligible mean vorticity, Haller *et al.* (2021) also derived TRA and $\overline{\text{TRA}}$, which measures trajectory rotation and successfully identifies elliptic LCS. Under these flow assumptions, $\overline{\text{TRA}}$ calculates the cumulative rotation of material streamline tangent vectors. For highly unsteady, or strongly rotational flows this approach is insufficient as streamlines no longer resemble material lines. Instead, we adopt the average rotation speed of material tangent vectors for relative Lagrangian velocities.

For a given relative Lagrangian velocity vector \mathbf{v}_d , we can define the unit vector pointing in that direction:

$$\mathbf{e}(t) = \frac{\mathbf{v}_d(t)}{|\mathbf{v}_d(t)|}. \tag{2.15}$$

Then,

$$r\alpha_{t_0}^{t_N} = \frac{1}{t_N - t_0} \int_{t_0}^{t_N} \left| \dot{\mathbf{e}}(t) - \frac{1}{2} \bar{\boldsymbol{\omega}}_{rb}(t) \times \mathbf{e}(t) \right| dt \tag{2.16}$$

is an objective measure of the average rotation speed of \mathbf{v}_d from time t_0 to t_N (Haller *et al.* 2021). The instantaneous limit of (2.16) also exists ($\alpha_{t_0} = \lim_{t_N \rightarrow t_0} \alpha_{t_0}^{t_N}$) and provides the instantaneous rate of rotation of the relative Lagrangian velocity vector with respect to the spatially averaged rotation. Furthermore, for a generic parameterised curve γ in three dimensions with a unit tangent velocity vector \mathbf{e} , the curvature of γ is $\dot{\mathbf{e}}(t)$. Thus, $r\alpha$ is intrinsically related to classical notions of geometry of flow structures as a form of integrated curvature that accounts for rotation of the surrounding fluid as well.

2.5. Extremely sparse sampling

As we show in the following sections, calculating $\bar{\mathbf{v}}(t)$ and $\bar{\boldsymbol{\omega}}(t)$ directly from trajectory data can become problematic at extremely sparse trajectory concentrations. By extremely sparse, we refer to trajectory concentrations less than $1\ell^{-2}$ and $1\ell^{-3}$, where ℓ is a characteristic length of structures in the flow (e.g. eddy length scale, obstruction dimensions). At this lower limit of trajectory concentration, we are fundamentally undersampling the prominent structures in the flow, and an accurate time-resolved depiction of spatially averaged flow behaviour is not reasonably expected.

That is, $\bar{\boldsymbol{\omega}}_{rb}(t)$ and $\bar{\mathbf{v}}(t)$ may start to display strong oscillations in time that are not representative of the underlying mean flow properties. These oscillations contribute unphysical fluctuations in \mathbf{v}_d and hinder our interpretation of relative stretching and rotation diagnostics. Oscillations can also be easily created by a small number of trajectories leaving the observation domain when only a small total are being measured. An example of this is shown in detail in § 3.1.2.

For these sparse situations, we suggest replacing the time-dependent bulk values $\bar{\boldsymbol{\omega}}_{rb}$, $\bar{\mathbf{v}}$, $\bar{\mathbf{x}}$ from ((2.8)–(2.11)) with their temporal averages over the integration window as

follows:

$$\mathbf{v}_{ref} = \frac{1}{t_1 - t_0} \int_{t_0}^{t_1} \bar{\mathbf{v}}(t) dt, \tag{2.17}$$

$$\boldsymbol{\omega}_{ref} = \frac{1}{t_1 - t_0} \int_{t_0}^{t_1} \bar{\boldsymbol{\omega}}_{rb}(t) dt, \tag{2.18}$$

$$\mathbf{x}_{ref} = \frac{1}{t_1 - t_0} \int_{t_0}^{t_1} \bar{\mathbf{x}}(t) dt. \tag{2.19}$$

While potentially oversimplifying non-stationary oscillations that naturally occur in the flow, using these time-averages provides a smooth estimate of total bulk motion for the time windows in which we are interested. We can then write the Lagrangian velocity relative to this prescribed reference frame as

$$\mathbf{v}_{d,ref}(\mathbf{x}(t)) = \mathbf{v}(\mathbf{x}(t)) - \mathbf{v}_{ref}(t) - \boldsymbol{\omega}_{ref}(t) \times (\mathbf{x} - \mathbf{x}_{ref}), \tag{2.20}$$

where \mathbf{v}_{ref} and $\boldsymbol{\omega}_{ref}$ are steady approximations of the bulk translation and rotation in the flow that are informed directly from flow measurements. This extremely sparse data handling is tested for multiple flows in the following sections, as is its ability to recreate accurate full-resolution α values. Since we are explicitly keeping track of a given choice of reference frame in (2.20), $\mathbf{v}_{d,ref}$, $\overline{\text{rTSE}}$ and α remain objective and indifferent to Euclidean frame changes.

If additional sources of velocity measurements are available, such as those from anemometers, remotely sensed ocean currents or wind tunnel probes, one may also estimate \mathbf{v}_{ref} and $\boldsymbol{\omega}_{ref}$ from those data, but this approach is not tested herein and should only be considered with caution. Using additional data sources to estimate bulk rigid-body rotation and translation is subject to errors, such as local fluctuations or imperfect placement in the flow, that are not discussed in the present work.

If one imposes $\mathbf{v}_{ref} = 0$, and $\boldsymbol{\omega}_{ref} = 0$, we can see that TRA and TSE are actually special cases of α and $\overline{\text{rTSE}}$. The present work, therefore, seeks to extend the prior success found by Haller *et al.* (2021) and Encinas-Bartos *et al.* (2022) with TRA and TSE to unsteady, time-varying flows. Further, the generalised relative Lagrangian velocity frameworks allows researchers to transparently maintain objectivity if a given reference frame is preferred.

3. Results

We now provide examples of α and $\overline{\text{rTSE}}$ analysis for numerical simulations and experimental observations. We evaluate our newly proposed metrics against an array of diagnostics suitable for different experimental flows, with each example detailing a specific advantage of α and/or $\overline{\text{rTSE}}$ over prior approaches.

3.1. Numerical simulations

3.1.1. Unsteady Bickley jet

For our first example we consider the unsteady Bickley jet. This is a 2-D geophysical model of a quasi-periodic zonal jet with adjacent migrating eddies whose Lagrangian dynamics have been studied in depth (see, e.g. Del-Castillo-Negrete & Morrison 1993;

U	L	k_n	ϵ_1	ϵ_2	ϵ_3	c_3	c_2	c_1
62.66 m s^{-1}	1770 km	$2n/r_0$	0.0075	0.15	0.3	$0.205U$	$0.461U$	$c_3 + \frac{k_2(\sqrt{5}-1)}{2k_1}(c_2 - c_3)$

Table 1. Bickley jet model parameters.




Figure 1. Evaluation of $r\alpha$ and $r\overline{\text{TSE}}$ (a,d) with material rotation (LAVD) and stretching (FTLE) (b,e), and traditional quasi-objective single trajectory diagnostics $\overline{\text{TRA}}$ and $\overline{\text{TSE}}$ (c,f), in the unsteady Bickley jet. Note the distinct central jet with rotating eddies on either side. Here $\overline{\text{TRA}}$ and $\overline{\text{TSE}}$ have difficulty accurately capturing material behaviour (stretching and rotation) due to the large bulk advection in the flow and temporal variability in the velocity field.

Rypina *et al.* 2007). This flow was also used as a benchmark for previous LCS comparisons (Hadjighasem *et al.* 2017). Its time-dependent stream function is given by

$$\psi(x, y, t) = -UL \tanh\left(\frac{y}{L}\right) + UL \operatorname{sech}^2\left(\frac{y}{L}\right) \operatorname{Re} \left[\sum_{n=1}^3 \epsilon_n \exp(-ik_n c_n t) \exp(ik_n x) \right]. \quad (3.1)$$

We use the parameters from Rypina *et al.* (2007) found in table 1.

In figure 1 we compare $r\alpha$ and $r\overline{\text{TSE}}$ with benchmarks of material rotation and stretching, the LAVD (Haller *et al.* 2016) and the FTLE (Haller 2015) on a dense grid of trajectories ($3600\ell^{-2}$, $\ell = \pi$), for an integration time of 30 days. This corresponds with approximately six full rotations and significant translation for the advected eddies. We find that $r\alpha$ is able to match the prominent features of the LAVD field by effectively identifying eddies boundaries as closed convex contours and separating them from the central jet. The $r\overline{\text{TSE}}$ ridges identify the edges of the central jet and edges of vortices as regions of significant stretching, similar to FTLE. Deviations primarily exist in the centres of the vortex cores where $\mathbf{v}_d(\mathbf{x}(t))$ does not evolve in a material way.

Previous work has shown how $\overline{\text{TRA}}$ and $\overline{\text{TSE}}$ can effectively reproduce the same LCS as LAVD and FTLE in slowly varying flows (Haller *et al.* 2021). In figure 1, we reveal how necessary this slowly varying assumption truly is. With the parameters in table 1, the Bickley jet is highly unsteady. For meaningful quasi-objective calculations, one requires $|\mathbf{v}_t|/|\dot{\mathbf{x}}| \ll 1$, but only 41 % of trajectories experience $|\mathbf{v}_t|/|\dot{\mathbf{x}}| < 1$ and only 0.03 % of particle paths show $|\mathbf{v}_t|/|\dot{\mathbf{x}}| < 0.1$. This suggests that the coherent structures in this flow are likely evolving or travelling much faster than a Lagrangian particle is able to trace




Figure 2. Comparison of relative rotation ($r\alpha$) and stretching ($r\overline{TSE}$) with non-objective trajectory length and initial particle speed diagnostics. Each diagnostic was progressively and randomly downsampled with the full-resolution product created using radial basis functions. For each panel, we include the outermost closed convex LAVD contours to highlight relative agreement of each sparse diagnostic with ground-truth material rotation.

them out. In fact, \overline{TRA} designates the eddies in figure 1 as exhibiting relatively low rotation. In addition, \overline{TSE} also suggests nearly everything outside the central jet is undergoing significant stretching, in contrast to the thin hyperbolic regions seen as FTLE ridges. For this degree of unsteadiness and strong eddy advection, $r\alpha$ and $r\overline{TSE}$ are more suitable than \overline{TRA} and \overline{TSE} for quantifying fluid rotation and stretching while requiring much less-structured data than LAVD and FTLE.

We now progressively and randomly downsample the number of trajectories to test the robustness of our objective diagnostics at low trajectory densities. In figure 2, we calculate $r\alpha$ and $r\overline{TSE}$ as well as the particle trajectory length (Mancho *et al.* 2013) and initial particle speed $|\mathbf{v}(t_0)|$. Although trajectory length and particle speed are not objective, they are either easily computed or commonly used diagnostics for experimental LPT studies (see, e.g. Fu, Biwole & Mathis 2015; Tauro, Piscopia & Grimaldi 2017; Rosi & Rival 2018). After calculating each diagnostic, we reconstruct the full-resolution diagnostic field using radial basis functions, similar to Encinas-Bartos *et al.* (2022). In addition, we include the outermost closed convex LAVD contours from figure 1 as a visual reference to aid in the comparisons.

At a trajectory concentration of $360\ell^{-2}$, $r\alpha$ looks relatively similar to its full resolution sampling, but with added noise. Five concentrations of high $r\alpha$ are still distinguishable from the surrounding flow. The thin hyperbolic structures identified by $r\overline{TSE}$ in the Bickley jet, however, are already beginning to disappear. This degradation of hyperbolic coherent structures $r\overline{TSE}$ is similar to the findings of Haller *et al.* (2021) for \overline{TSE} and Mowlavi *et al.* (2022) for \overline{TSE} and FTLE. The trajectory length reveals fluid particles in the central jet have the longest trajectories over this time window when compared with nearby particles in the adjacent eddies. The initial velocity magnitude $|\mathbf{v}|$ also appears largest in the centre of flow domain, and suggests some oscillating feature may be present.

Dropping the trajectory resolution by a factor of 100 to $3.5\ell^{-2}$, we again reconstruct the original diagnostic resolution and include the location of trajectory positions with white crosses. At this concentration, $r\alpha$ still suggests five distinct local rotation maxima corresponding to the locations of the five eddies. The interpolated $r\overline{TSE}$ field no longer recreates any of the features seen at full resolution and hence provides no meaningful information. The trajectory length is also contradicting features that were revealed with 100 times the number of trajectories. At this resolution, we are testing both the ability

Relative fluid stretching and rotation for sparse trajectories

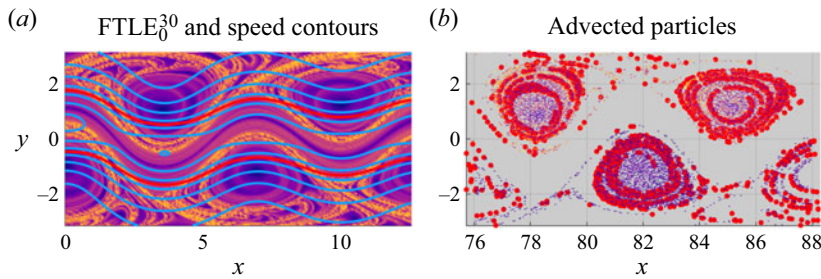


Figure 3. Comparison of particle speed, FTLE and the evolution of fluid particles advected from inside a given velocity level set.

of the metric to provide useful flow descriptors, as well as the interpolation scheme. We believe, however, that there is limited influence from the interpolation scheme because $|\mathbf{v}|$ is practically constant across all resolutions. We further test the impact of interpolation schemes in the Appendix, and refer the reader to Encinas-Bartos *et al.* (2022) for a deeper investigation of those schemes in a similar context. This trajectory concentration is well below the limits of the previous sparse trajectory studies of Lekien & Ross (2010) and Mowlavi *et al.* (2022).

From this initial investigation it appears that $|\mathbf{v}|$ may be the best option for identifying structures with low-resolution data. The dominant flow feature suggested by $|\mathbf{v}|$ is, in fact, misleading, even at full resolution. Figure 3(a) shows an overlay of $|\mathbf{v}|$ contours with the FTLE field in greyscale. The distribution of $|\mathbf{v}|$ suggests two distinct peaks with a distribution minimum, s_δ , separating the central and outer regions. The fast moving central core is highlighted as the region inside the red $|\mathbf{v}| = s_\delta$ contour in figure 3(a). Examination of figure 3(a) show that the jet and eddy structures suggested by the FTLE field have boundaries that are actually transverse to all of the velocity contours.

In figure 3(b) we plot the final position of fluid particles from *only* the velocity core ($|\mathbf{v}| \geq s_\delta$), coloured by their $|\mathbf{v}(t_0)|$ value, after 30 days of advection. Particles corresponding to the red boundary have been advected as well. It is clear that the velocity-based jet is not actually a coherent structure as the proposed feature has been stretched inside and around the eddies, as well as down the eastward jet, with initial particle speed giving no indication of what feature a particle should be attached to.

For comparison, we show the evolution of the α structures, those that were similarly resilient to downsampling. In figure 4(a) we plot α , with rotationally coherent structure boundaries identified as the outermost closed convex contours of α , as has been previously used in other elliptic LCS methods (Haller *et al.* 2016; Encinas-Bartos *et al.* 2022). Figure 4(c) shows the location of the same α -coloured fluid particles after 30 days of advection. The α eddy boundaries have minimally deformed after being advected downstream more than 20 times their diameter downstream. This time also corresponds with approximately six full rotations of the eddy, ample time for inaccurate eddy boundaries to experience filamentation. This strong coherence can be attributed to the mathematical definition of α as a measure of relative fluid rotation and suggest we have effectively and objectively identified the true rotationally coherent structures in our flow field.

The ability of α to identify these same eddies at a random lower data concentration ($3.5\ell^{-2}$) is tested in figure 4(b,d). We again identify eddies as the outermost closed convex contours of α at $t = 0$ (figure 4b) and advect them for 30 days. Comparing figures 4(a) and 4(b), we see the leftmost eddy boundary extending into the central jet and differences

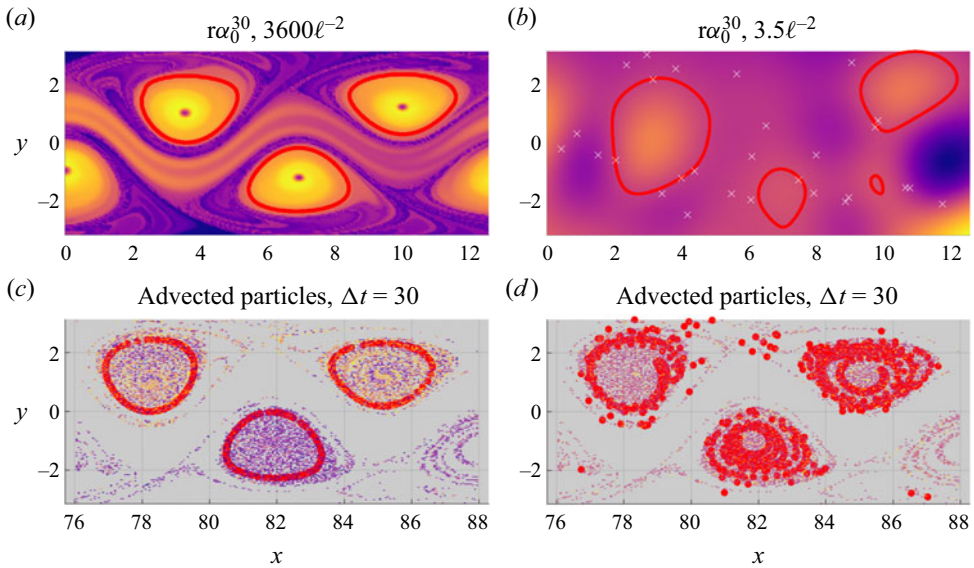


Figure 4. Outermost closed convex $r\alpha$ level-set contours overlaid on the $r\alpha$ field. Advected position of $r\alpha$ contours and surrounding fluid, coloured by initial $r\alpha$ values.

arising in the exact eddy boundary location. Testing these boundaries under advection, we find the red boundary particles have remain primarily inside the eddies we sought to identify (figure 4d), but with a much higher degree of filamentation.

The ability to identify structure boundaries is as much of a test of the interpolation scheme used as of $r\alpha$ to maintain meaningful values at lower resolution. When there are fewer than one trajectory per coherent structure, we will inevitably misrepresent some flow features in interpolated fields and $r\alpha$ and $r\overline{TSE}$ are more reliably interpreted as point measurements of local behaviour. As mentioned previously, we investigate interpolation scheme impacts in the Appendix. Here, we investigate the ability of low-concentration data to provide accurate reference frame values $\bar{\mathbf{v}}$, $\bar{\boldsymbol{\omega}}_e$ and $\bar{\mathbf{x}}$ and meaningful pointwise $r\alpha$, independent of the interpolation scheme, with the following experiment.

Starting with full-resolution ($3600\ell^{-2}$) Bickley jet trajectories, we first calculate ground-truth values for $r\alpha(\mathbf{x}_0)_{full}$. We then randomly subsample the full resolution trajectory 10^5 times at a given concentration. For each selection of trajectories, we recalculate $r\alpha(\mathbf{x}_0)_{low}$ using the lower-resolution data, and calculate the correlation coefficient (R^2) of $r\alpha(\mathbf{x}_0)_{low}$ and $r\alpha(\mathbf{x}_0)_{full}$ at the corresponding locations. We then obtain a distribution of R^2 values that details how closely we can approximate the full-resolution measurements for a large number of potential trajectory orientations. We then repeat this process at progressively lower concentrations and witness the rate at which the accuracy degrades.

The findings from this experiment are summarised in table 2. The mean value of R^2 slowly decreases with increasing sparsity but stays above 0.99, with a standard deviation ≤ 0.008 down to $16\ell^{-2}$. The mean of R^2 reduces more dramatically around the extremely sparse sampling threshold of $1\ell^{-2}$. This is precisely the trajectory concentration at which we suggest using time-averaged \mathbf{v}_{ref} , $\boldsymbol{\omega}_{ref}$ and \mathbf{x}_{ref} values. The impact of this choice is seen in the final column where mean R^2 values nearly return to the $4\ell^{-2}$ value. This suggests a

ℓ^{-2} :	128	64	32	16	8	4	2	1	1 (ref)
μ :	0.997	0.997	0.996	0.991	0.98	0.95	0.87	0.68	0.93
σ :	2×10^{-4}	0.001	0.003	0.008	0.018	0.045	0.11	0.25	0.14

Table 2. Bickley jet downsampling effect on α accuracy. Mean (μ) and standard deviation (σ) of R^2 values correlating downsampled α with full-resolution calculations. Raw α values maintain a strong correlation until the extremely sparse sampling threshold ($1\ell^{-2}$), at which point steady time-averaged reference values (‘ref’, § 2.5) allow for a significant improvement in accuracy again.

significant improvement in pointwise α accuracy, and meaningful rotation diagnostics for extremely sparse data.

3.1.2. AVISO

We now consider the identification of elliptic Lagrangian flow structures from 2-D ocean satellite altimetry data provided by AVISO which has been the focus of several coherent structure studies (see, e.g. Haller *et al.* 2016, 2021). The zonal and meridional component $\mathbf{v} = (v_1, v_2)$ of the ocean currents are derived from the sea-surface height profile

$$fv_2(\mathbf{x}, t) = \frac{1}{\rho} \frac{\partial}{\partial x} p(\mathbf{x}, t), \tag{3.2}$$

$$fv_1(\mathbf{x}, t) = \frac{1}{\rho} \frac{\partial}{\partial y} p(\mathbf{x}, t), \tag{3.3}$$

$$-g = \frac{1}{\rho} \frac{\partial}{\partial z} p(\mathbf{x}, t), \tag{3.4}$$

where $p(\mathbf{x}, t)$ is the pressure, ρ is the fluid density, f is the Coriolis parameter and g is the Earth’s acceleration. The daily-gridded velocity data are freely available from the Copernicus Marine Environment Monitoring Service. Although this is an observational product, the processed nature of the product prevents many of the complications that we will see with our experimental examples in the following section. Our analysis focuses on the North Atlantic Gulf Stream between longitudes 70°W and 55°W and latitudes 30°N and 45°N , spanning September and October 2006. We start with an initial grid of trajectories consisting of 150×150 points. Using a characteristic length scale of mesoscale eddies ($\ell = 100$ km), this density roughly corresponds to $100\ell^{-2}$.

Elsewhere in the Atlantic, Haller *et al.* (2021) and Encinas-Bartos *et al.* (2022) have shown that the AVISO velocity field is slowly varying, with a vanishing spatially averaged vorticity. Furthermore, these authors showed that $\overline{\text{TRA}}$ effectively identifies ocean eddies, and outperforms other sparse trajectory rotation diagnostics. As such, in figure 5 we compare the α (figure 5a) with $\overline{\text{TRA}}$ (figure 5b), using LAVD as the ground truth of material rotation (figure 5c). The white dashed box highlights the meanders of the Gulf Stream. Three eddy boundaries are identified as outmost closed convex contours of LAVD in figure 5(a–c).

When compared with LAVD, $\overline{\text{TRA}}$ faithfully approximates the LAVD field, identifying the meanders of the Gulf Stream, high rotation in the adjacent eddies and relatively low rotation away from these central features. At this trajectory density, figure 5(a) reveals that α can match the ability of $\overline{\text{TRA}}$ to visualise the correct flow structure topology.




Figure 5. Overview of rotation diagnostics for the AVISO data in the Gulf Stream at full resolution ($100\ell^{-2}$). The white contours are the outermost convex contours of the LAVD field. The dashed white box highlights the meanders of the Gulf Stream.

This stems from the similarity in $\overline{\text{TRA}}$ and α definitions for small bulk advection and rotation, as highlighted in § 2.5.

We compare the sparse data performance of α against $\overline{\text{TRA}}$ in figure 6 by progressively and randomly subsampling using the same methodology as in § 3.1.1. We again interpolate back to full resolution using a radial basis function for each successive subsampling. We also include the major eddy boundaries from figure 5.

At a trajectory concentration of $10\ell^{-2}$, $\overline{\text{TRA}}$ (figure 6a) and α (figure 6d) appear qualitatively similar, with both suggesting relatively high rotation in the Gulf Stream meanders, and three distinct elliptical rotational maxima aligned with the eddy boundaries. For a trajectory density of $1\ell^{-2}$, α (figure 6e) is qualitatively indistinguishable from $\overline{\text{TRA}}$ (figure 6b) with two of the three mesoscale LAVD eddies still visible as local maxima in both diagnostic fields. For this region of the ocean, we often have ocean drifter observations [e.g. Global Drifter Program (GDP)] at concentrations between $1\ell^{-2}$ and $0.1\ell^{-2}$, as discussed in § 3.2.1.

As we pass the extremely sparse sampling threshold to $0.1\ell^{-2}$, the leftmost eddy is still identifiable as a $\overline{\text{TRA}}$ and α maximum (figures 6c and 6f, respectively) due to the presence of the single trajectory originating inside its boundary (white cross). The remaining eddies are no longer visible for either metric due to a lack of data, but $\overline{\text{TRA}}$ still highlights a relatively higher rotation in the Gulf Stream. The local maximum $\overline{\text{TRA}}$ within the dashed white box may be classified incorrectly as an eddy, but it actually corresponds to a region of high material rotation induced by shear in the Gulf Stream. In contrast, α has created multiple local maxima near the coast of Nova Scotia not present at higher concentrations.

At this extremely sparse resolution, we again replace $\bar{\omega}(t)$, $\bar{\mathbf{v}}(t)$, and $\bar{\mathbf{x}}(t)$ with their time-averaged steady approximations, \mathbf{v}_{ref} , $\boldsymbol{\omega}_{ref}$, \mathbf{x}_{ref} from (2.17)–(2.19). As with the Bickley jet statistics, this improves the accuracy of α with α_{ref} displayed in figure 6(i). With the steady reference frame, α_{ref} and $\overline{\text{TRA}}$ again show nearly indistinguishable flow topologies. The coastal local maxima disappear, and the local maximum appears in the shear dominated meanders of the Gulf Stream. Furthermore, if we apply this steady reference frame change to higher concentrations (figure 6g and h), we obtain flow structure visualisations comparable to α . This suggests that obtaining meaningful results using the extremely sparse steady reference-frame approach is resilient over a range of concentrations.

Relative fluid stretching and rotation for sparse trajectories




Figure 6. Comparison for the AVISO data between objective rotational sparse data diagnostic α , α_{ref} and its quasi-objective single-trajectory counterpart TRA under varying trajectory densities ($10\ell^{-2}$, $1\ell^{-2}$, $0.1\ell^{-2}$). The white markers denote the trajectory endpoints.

This effect stems from the effect of undersampling on bulk rotation and translation estimates. In figure 7, we display the $\bar{\omega}(t)$ and $\bar{v}(t)$ estimates for each trajectory concentration. At $100\ell^{-2}$ the mean flow properties are nearly constant over 15 days. This is consistent with the expectation from earlier studies that Eulerian features do not change significantly over such relatively small time scales in the ocean. As we progressively reduce the trajectory density, however, $\bar{\omega}(t)$ and $\bar{v}(t)$ start displaying strong oscillations in time. At extremely low trajectory densities ($< 1\ell^{-2}$), the spatial averages are no longer representative of the underlying mean flow properties (dashed red curves in figure 7). These oscillations contribute unphysical fluctuations in the direction of v_d and hinder our ability to measure rotation with α .

To further quantify the impact of sparse sampling on α , we repeat the R^2 distribution analysis that was performed for the unsteady Bickley jet. Using AVISO surface current data to generate an initial baseline at $100\ell^{-2}$, we run 10^4 random subsamplings at a range of concentration levels. The summary statistics of this experiment can be found in table 3. A consistent decrease in the R^2 mean, and increase in R^2 standard deviation is found as the data concentration levels decrease, though the level of α accuracy remains quite high. Even at a low concentration level of $1\ell^{-2}$, we obtain a mean R^2 value of 0.92.




Figure 7. Data-based estimates for bulk translation velocity $|\bar{v}(t)|$ and rotation $\bar{\omega}(t)$ under varying trajectory density ($100\ell^{-2}$, $10\ell^{-2}$, $1\ell^{-2}$, $0.1\ell^{-2}$).

ℓ^{-2} :	64	32	16	8	4	2	1	0.1	1 (ref)	0.1 (ref)
μ :	0.999	0.998	0.996	0.992	0.98	0.96	0.92	0.45	0.95	0.89
σ :	2×10^{-4}	7×10^{-4}	0.002	0.003	0.008	0.02	0.04	0.21	0.02	0.13

Table 3. Downsampling effect on $r\alpha$ accuracy for AVISO Gulf Stream surface currents. Mean and standard deviation of R^2 values relating downsampled $r\alpha$ with full-resolution calculations for 10^4 random subsamplings at each resolution. Raw $r\alpha$ values maintain a strong correlation until the extremely sparse threshold ($1\ell^{-2}$), at which point time-averaged reference values (ref) allow for a significant improvement in accuracy again.

The most precipitous drop in $r\alpha$ accuracy is found from $1\ell^{-2}$ to $0.1\ell^{-2}$. One such example of this poor performance was highlighted in figure 6. We further validate the improvement in $r\alpha$ accuracy at this extreme sparsity level in the last two columns of table 3. After transitioning to time-averaged reference frame values, there is a significant increase in accuracy, with the mean of R^2 jumping from 0.92 to 0.95 for $1\ell^{-2}$, and from 0.45 to 0.89 for $0.1\ell^{-2}$. With this simple reference-frame methodology we find $r\alpha$ is a promising Lagrangian flow rotation diagnostic for extremely sparse data.

3.1.3. Stationary concentrated vortex

Our third example considers the identification of elliptic coherent structures inside the stationary concentrated vortex model (SCVM) of Onishchenko *et al.* (2021). The SCVM represents finite-size cyclones in the Earth’s atmosphere as steady axially symmetric solutions to the Euler equations. Both the vertical and radial extent of the vortex is controlled by model parameters. The SCVM fundamentally differs from the previous flows as it is a highly rotational flow with no bulk translation. Our investigation of the SCVM helps confirm that $r\alpha$ and $r\sqrt{\text{TSE}}$ are robust diagnostics for approximating the shape and location of LCS and are not tailored to a specific kind of flow experiment (e.g. both high and low swirl number; Gupta, Lilley & Syred 1984).

The SCVM consists of inner and outer regions, an internal upwards motion and external downwards motion. A distinct central torus was also identified by Aksamit (2023) as a

LCS. The cyclone is further divided vertically into a centripetal and centrifugal flow with upward-moving fluid recirculated from the top of the vortex back down.

The SCVM was originally derived as an axisymmetric flow from a stream function $\Psi(r, v)$ and can be written in cylindrical coordinates as

$$v_r = -v_0 \frac{r}{L} \left(1 - \frac{z}{L}\right) \exp\left(-\frac{z}{L} - \frac{r^2}{r_0^2}\right), \quad (3.5)$$

$$v_\theta = \pm v_{\theta_0} \frac{z}{L} \exp\left(-\frac{z}{L} - \frac{r^2}{r_0^2}\right), \quad (3.6)$$

$$v_z = 2v_0 \frac{z}{L} \left(1 - \frac{r^2}{r_0^2}\right) \exp\left(-\frac{z}{L} - \frac{r^2}{r_0^2}\right), \quad (3.7)$$

where r_0 is the radial extent separating inner and outer vortex structures, L is the height of maximum vertical velocity at $r = 0$ and v_0 and v_{θ_0} are characteristic velocities. For our example, we use $v_0 = 1$, $L = 10$, $r_0 = 1$ and $v_{\theta_0} = 9$. We consider our characteristic length of the flow to be equal to the separation of inner and outer motions, $\ell = r_0 = 1$, and the approximate radius of the main central torus.

In [figure 8](#) we compare the objective rotational diagnostics $r\alpha$ and LAVD advected for $\Delta t = 500$ with the vorticity magnitude, $|\omega|$ and the widely used lambda-2 criterion of Jeong & Hussain (1995). If \mathbf{J} is the velocity gradient tensor, and \mathbf{S} and $\mathbf{\Omega}$ are the symmetric and antisymmetric parts, the lambda-2 criterion of Jeong & Hussain (1995) suggests that regions with negative intermediate eigenvalues of $\mathbf{S}^2 + \mathbf{\Omega}^2$ (labelled λ_2) are part of a vortex core. Stationarity of the flow suggests $\overline{\text{TRA}}$ would also be a suitable comparison, but in light of the strong performance of $r\alpha$ against $\overline{\text{TRA}}$ in the previous two examples, we have left it out for compactness of presentation.

From initial high spatial resolution ($2000\ell^{-3}$) calculations, a 2-D slice along the $y = 0$ plane is shown for each field in [figure 8](#). We overlay these plots with the level-set contours of the SCVM stream function $\Psi(r, v)$. As this is a stationary axisymmetric flow, Ψ -contours provide us a ground truth of transport barrier structure. A qualitative contour comparison of the sparse metrics ($r\alpha$ and $r\overline{\text{TSE}}$) with Ψ show a comparable level of structure agreement as one can find from LAVD and FTLE structures. We find that $r\alpha$ and LAVD contours suggest very similar rotational behaviour, with both indicating the same location of a central torus vortex surrounded by recirculating flow. Both $r\overline{\text{TSE}}$ and FTLE show strong spiraling into the centre of the vortices identified by $r\alpha$ and LAVD. One notable difference between sparse and spatial-gradient-reliant diagnostics exists in the centre of the toroidal vortex core in $r\overline{\text{TSE}}$ which again have unnaturally low values when compared with FTLE, as we also saw in the unsteady Bickley jet.

We see a stark difference, however, when comparing Ψ with $|\omega|$ and λ_2 fields. Contours of $|\omega|$ and λ_2 are often orthogonal to Ψ -contours, suggesting a disagreement in structure identification. We further investigate these qualitative findings by testing the ability of our rotation diagnostics to identify the dominant elliptic coherent structures in this vortex flow. Potential vortex boundaries are identified as the set of outermost closed convex contours for $r\alpha$, LAVD and $|\omega|$. Following the derivation of the λ_2 criterion, we use $\lambda_2 < 0$ as the designation of strong rotational motion (Jeong & Hussain 1995). Fluid parcels from inside the boundaries of these features are then advected for twice the integration time used for the Lagrangian diagnostics, $\Delta t = 1000$. This approach allows us to test both the

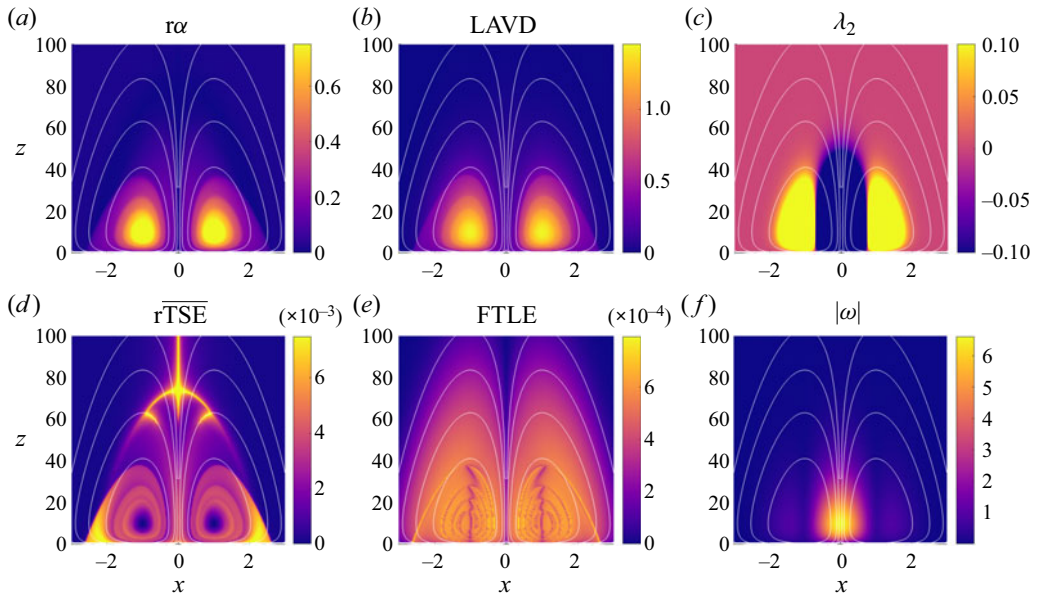


Figure 8. Comparison of rotation diagnostics $r\alpha$, LAVD, λ_2 and $|\omega|$ for the SCVM in the $y = 0$ plane. White contour lines are contours of the SCVM stream function Ψ , indicating circulation patterns in the model. Note the agreement between $r\alpha$, LAVD, $r\overline{TSE}$ and FTLE features with the stream function structures, and the dissimilarity of $|\omega|$ and λ_2 features.

Lagrangian and Eulerian diagnostics to identify features that maintain their coherence beyond the time horizon of data used for the initial boundary calculation.

These fluid particle paths are shown in [figure 9](#) with trajectories coloured by the initial diagnostic field value. We include the corresponding [figure 8](#) diagnostic field on the $y = 0$ plane as a reference. The vortex core identified by $r\alpha$ and LAVD maintain coherence over time, with seemingly no deviation of fluid from the domain identified as the central torus vortex ([figure 9a,b](#)). Given that $r\alpha$ does not rely on the same 3-D velocity gradient data as LAVD, $r\alpha$ shows exceptional performance in identifying this LCS. Fluid initially inside λ_2 and $|\omega|$ structures have instead maintained no coherence, and are sufficiently dispersed in the flow. The high $|\omega|$ ([figure 9d](#)) fluid parcels are now mixed throughout the entire inner SCVM ‘beehive’ domain, inhabiting regions with both high and low $|\omega|$ values. The $\lambda_2 < 0$ fluid is recirculating in both the inner and outer regions of the SCVM ([figure 9c](#)), with the λ_2 -estimated vortex core parcels inhabiting the domain of both negative and positive λ_2 values. This is distinct from LAVD and $r\alpha$ vortex trajectories that stay confined to the regions of high LAVD and $r\alpha$, respectively.

The ability of $r\alpha$ to identify these same LCS at increasing sparsity is shown in [figure 10](#). As our only 3-D numerical example, we utilise a random set of initial particle trajectories in three dimensions and extract isosurfaces from a reconstructed $r\alpha$ field using natural neighbour interpolation. We utilise the same $r\alpha$ value for 3-D level-set extraction as the outermost closed convex contour used to test vortex boundaries ([figure 9](#)). Plots of reconstructed $r\alpha$ fields on the $y = 0$ plane ([figure 10a–c](#)) reveal that small deviations from this range would not significantly change the topology of the surfaces in [figure 10\(d–f\)](#). Note the difference of z -value ranges between [figures 10\(a–c\)](#) and [10\(d–f\)](#).

In [figure 10](#), $r\alpha$ at $10\ell^{-3}$ looks only moderately different from $2000\ell^{-3}$. Two distinct intersections of the central torus with the $y = 0$ plane are still evident in white, and

Relative fluid stretching and rotation for sparse trajectories

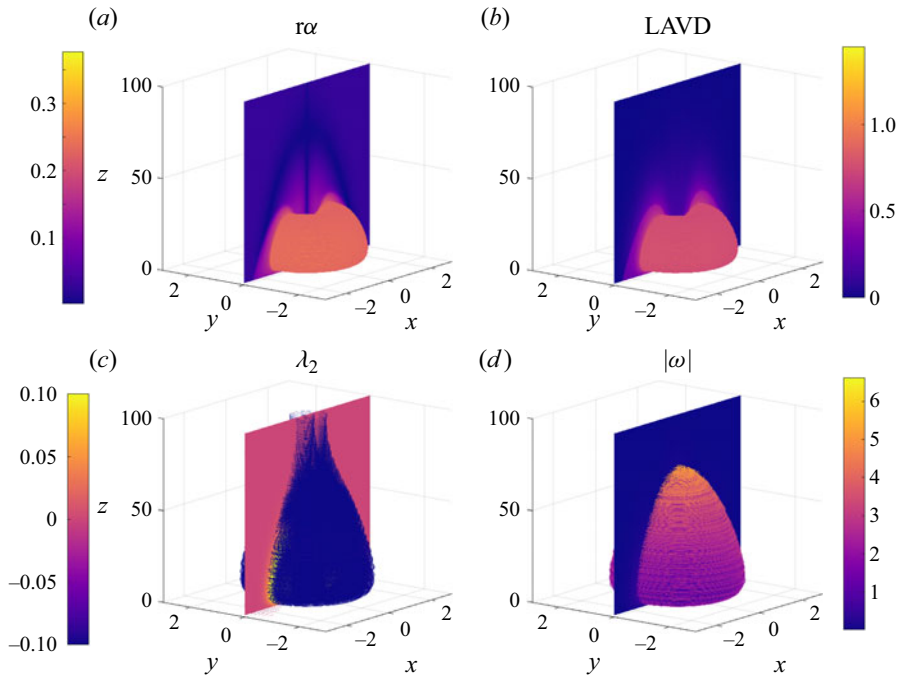


Figure 9. Trajectories corresponding to coherent structures identified by the four diagnostics, calculated for the same integration times.




Figure 10. (a–c) Values of $r\alpha_0^{500}$ on the $y = 0$ plane reconstructed from randomly selected trajectories at decreasing particle density. The inner vortex is extracted at each concentration as the closed contours corresponding to the level-set value from figure 8. (d–f) Comparison of the inner vortex surface in $r\alpha_0^{500}$ fields reconstructed from randomly selected trajectories at decreasing particle density. The inner vortex is extracted in each scenario as closed contours corresponding to the outermost closed convex contour in figure 9.

the annular shape is clearly visible in the associated isosurface. At $1\ell^{-3}$, the ‘doughnut hole’ above $x = 0$ becomes less obvious, but the horizontal and vertical extent of the central torus is still accurate. For $0.1\ell^{-3}$, we use ((2.17)–(2.18)) to implement the same bulk time-averaged reference values for \mathbf{v}_{ref} , $\boldsymbol{\omega}_{ref}$, \mathbf{x}_{ref} . At this extremely sparse sampling, the correct extent of strong central rotation is still evident, as is a suggestion of the inner beehive feature, but the vortex surface is now sufficiently deformed that it no longer resembles a torus.

3.2. Experimental data

In the previous section we investigated the ability of $\overline{\alpha}$ and $\overline{\text{rTSE}}$ to accurately represent elliptic and hyperbolic LCS at increasing sparsity in numerical data. In this section, we show how the sparse capabilities of $\overline{\alpha}$ apply to real experimental data for a 2-D and 3-D flow.

3.2.1. GDP database

For our first experimental data set we consider a set of drifters from the GDP (Lumpkin & Centurioni 2019) located in the North Atlantic Ocean. The GDP data set contains more than 40 000 drifters spanning the last four decades. Roughly 1200 drifters are currently active worldwide and report their position every 6 hours (Lumpkin & Pazos 2007).

Specifically, we focus here on a subset of drifters active over 15 days from 20 September to 4 October 2006 in the Gulf Stream (figure 11) during the same time period of analysis as in § 3.1.2. In figure 11, we reconstruct diagnostic fields from the true GDP tracks with a drifter density of $0.63\ell^{-2}$, for $\ell = 100$ km. We are thus in an extremely sparse data setting ($< 1\ell^{-2}$) and we again compute $\overline{\alpha}$ by replacing the time-dependent mean flow properties with time-averaged $\boldsymbol{\omega}_{ref}$, \mathbf{v}_{ref} and \mathbf{x}_{ref} from the drifter data.

Following the sparse eddy identification study by Encinas-Bartos *et al.* (2022), we compare two sparse drifter-based rotation diagnostics, $\overline{\alpha}_{ref}$ in figure 11(a) and $\overline{\text{TRA}}$ in figure 11(b), with commonly used trajectory-averaged kinetic energy (KE; figure 11(c) (e.g. Lumpkin 2016), and the full-resolution LAVD field using AVISO data from § 3.1.2. Though the AVISO current data may not be representative of the full spatial scale of flow features affecting the drifters, we include LAVD as a mesoscale comparison for general flow behaviour.

We do not compare with a direct ‘looping’ count as a user-defined threshold is inevitably required and differentiating between looping and non-looping trajectories is user-dependent and non-trivial (Encinas-Bartos *et al.* 2022). Instead, we extract looping segments for the reader to see from the drifter trajectories as proposed by the frame-dependent algorithm from Lumpkin (2016). The white trajectory segments indicate cyclonic loopers which frequently arise in connection with ocean eddies (see, e.g. Griffa, Lumpkin & Veneziani 2008; Dong *et al.* 2011). The full black curves in figure 11(a,b,d) denote the approximate vortex boundaries respectively obtained from drifter-based $\overline{\text{TRA}}$ and $\overline{\alpha}_{ref}$, and AVISO-based LAVD, respectively.

Both $\overline{\alpha}_{ref}$ and $\overline{\text{TRA}}$ are able to detect rotational features just to the south of the Gulf Stream from the extremely sparse GDP data. Cyclonic loopers are also observed in the regions surrounding the principal local maxima of $\overline{\alpha}_{ref}$ and $\overline{\text{TRA}}$. The leftmost eddy contains two looping drifters that display high rotational values (figure 11a,b). In contrast, $\overline{\text{KE}}$ fails to clearly identify the leftmost eddy since it displays sharp gradients between the two drifters inside the eddy (figure 11c). Both $\overline{\alpha}_{ref}$ and $\overline{\text{TRA}}$ display a weaker rotational

Relative fluid stretching and rotation for sparse trajectories

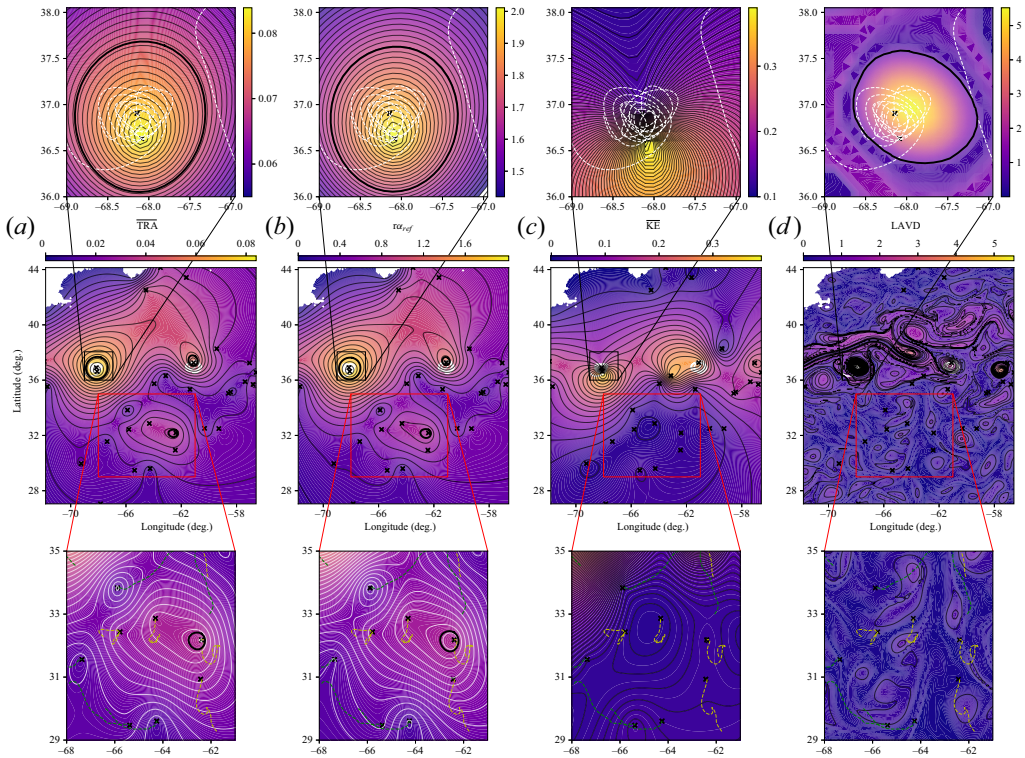


Figure 11. Drifter-based $\overline{\text{TRA}}$ (a), α (b), $\overline{\text{KE}}$ (c) and LAVD (d) in the North Atlantic Gulf Stream. The black crosses indicate drifter positions.

feature at roughly 32°N latitude and 64°W longitude. No loopers are identified in this region, but the yellow trajectories in the high α and $\overline{\text{TRA}}$ zones show greater circulating behaviour with respect to the surrounding green trajectories (figure 11, bottom insets), suggesting stronger rotational motion in the surface currents. There is no feature that resembles any vortical structure in the $\overline{\text{KE}}$ field in this region.

In this ocean buoy experiment, we were able to identify eddies adjacent to the Gulf Stream in both $\overline{\text{TRA}}$ and α fields that were not identified using existing sparse data diagnostics. In contrast with the $\overline{\text{TRA}}$, which requires additional assumptions on the frame of reference, the α returns structures that are valid for all observers. This suggests that α may be beneficial for systematic eddy analysis of the GDP data set that can complement remote-sensing based tools that lack submesoscale motion resolutions. In addition, α shows great promise as a buoy diagnostic for focused buoy deployments (e.g. LASER and GLAD) where buoy concentrations higher than $1\ell^{-2}$ will result in more refined descriptions of the flow.

3.2.2. Large-scale LPT experiment

Our last example involves a novel, large-volume 3-D LPT velocimetry data set from an industrial-scale wind tunnel (Hou *et al.* 2021). Utilising glare-point spacing on large soap bubbles, Hou *et al.* (2021) were able to investigate the vortical wake behind a tractor-trailer model at a 9° yaw angle. For details on the methodology, see Kaiser & Rival (2023). We

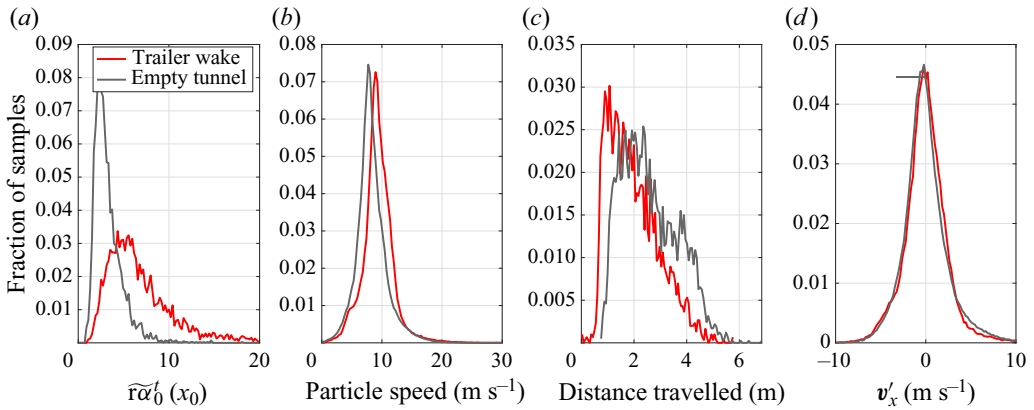


Figure 12. Comparison of discrete probability histograms of frame-dependent metrics for the truck wake vortex and object-free uniform flow experiments at $\sim 1\ell^{-3}$.

compare two wind-tunnel experiments, one with the tractor–trailer model obstructing the flow, and a uniform flow without the tractor–trailer.

Each experiment provides around 300 s of trajectory data, sampled at 150 Hz with an average windspeed of 10 m s^{-1} . The bubble generator and camera setup produce time-varying particle concentrations at the very lower end of our computation spectrum. For measurements of the tractor–trailer wake, trajectory concentrations oscillate between 0.2 and $1.5\ell^{-3}$ with a mean concentration of $0.6\ell^{-3}$. The uniform flow measurements have a slightly higher resolution, ranging from 0.5 to $3\ell^{-3}$ and a mean concentration of $1.3\ell^{-3}$. To avoid significant oscillations in relative Lagrangian velocities, we again utilise the time-averaged reference frame determined by the bubble trajectories. Furthermore, as we are using trajectory data that quickly transits through the observation volume, we investigate 10-frame increments and require all trajectories to be present for every frame. At these extremely low concentrations, a single particle that disappears and reappears during the integration window may have a substantial effect on the deformation velocity of other fluid particles. A 10-frame integration time was chosen as a balance between the length of integration time and the number of tracks available to calculate meaningful mean values.

At the low trajectory concentration available in this turbulence experiment, meaningful time-resolved spatial derivatives are unattainable. Instead of relying on a spatial interpolant to visualise strong regions of rotation, we now demonstrate a complementary statistical approach using this sparse data set. In [figure 12](#) we show the discrete probability histograms for common LPT flow diagnostics in addition to $\alpha\omega$ analysis. The distribution of the particle speed over the entire 300 s ([figure 12b](#)) shows a slight speed up of the flow for the trailer wake, but the bubble trajectories travel for a shorter distance on average ([figure 12c](#)). The fluctuating streamwise components of bubble velocities were also nearly identical for the two flows.

Sensor development for large-scale LPT is still an active field of research. Given the uncertainty in signal noise from nascent methodologies, such as the glare-point method used to generate for these data, flow diagnostics that maintain their utility for unknown noise levels can help to expedite the development of large-scale measurement tools. With the available information from the industrial-scale wind tunnel, none of the diagnostics in

Relative fluid stretching and rotation for sparse trajectories




Figure 13. (a,b) Binned median value of $r\alpha_0^t(x_0)$ after projecting in the streamwise direction onto the y - z plane. The trailer wake has high rotation in its core, and relatively higher $\bar{r}\alpha_0^t$ values distributed across the flow. The empty tunnel, in contrast has a low rotational core, with higher $\bar{r}\alpha_0^t$ appear only on the edges of the measurement domain. This is confirmed by binning the $r\alpha_0^t$ values by radial distance from the flow centre, and calculating median values ($\bar{r}\alpha_0^t(r)$, panel c). Note the radial centres and axes are marked by blue dots and lines in panels (a,b).

figure 12(b-d) present a clear picture of the wake dynamics. In fact, v'_x cannot distinguish between the two flows.

To complement these approaches, we compute 10-frame $r\alpha$ values beginning at every frame available. This gives us a distribution of $r\alpha_t^{t+10}(x_0)$ for a wide range of t and x_0 . From a set of trajectories at each time t , we then calculate the median of $r\alpha$ ($\bar{r}\alpha$) to quantify the general behaviour of the flow in each time window. The discrete probability histogram of $\bar{r}\alpha$ for the two flows is shown in figure 12(a). The empty tunnel data are concentrated around a low peak, suggesting a relatively uniform flow domain and an accumulation of small turbulent oscillations, but no major rotational features. In contrast, the trailer wake data are distinguished by a larger range of values, with a heavy tail towards considerable rotation and high peak $\bar{r}\alpha$. The physical meaning of this data suggests a more complex flow, with a higher degree of rotation, in this case because of the vortical wake. Even at this low concentration, and with this experimental methodology, we are able to identify distinguished rotating features in the flow.

We can begin to distinguish the spatial features of the flow by analysing the long-time trends in $r\alpha$. In figure 13 we project all initial particle locations to the y - z plane, and bin them in a 100×100 grid. We again calculate $\bar{r}\alpha$ for each bin and plot this in figure 13(a,b). The trailer wake in figure 13(a) has high rotation near its centre, and relatively higher $r\alpha$ distributed throughout the flow domain when compared with the quiescent centre and lower values in figure 13(b). We quantify the spatial dependence of rotation on flow geometry by binning $r\alpha$ in 10 bins along the radial distance from the centre of the flow. We show this radial axis as a dot and line in figure 13(a,b). Again, the median $r\alpha$ values from these bins shows the trailer wake generates a much greater degree of rotation on the fluid, starting from a minimum where little data are available, with a local maximum in the vortex core, then gradually increasing as we move away in the flow. At nearly every location, the empty tunnel generates limited fluid rotation with lower $r\alpha$ values.

The results here provide a good baseline for experimental LPT data that were not designed with $r\alpha$ analysis in mind. As future experiments may be designed that consider $r\alpha$ requirements, further improvement on the identification and analysis of coherent rotational flow features is expected. Furthermore, given the connections between $r\alpha$, curvature (e.g. § 2.4) and rotation times, concerted future effort may help connect classical notions of

geometry, eddy turnover times, spectral analysis and sparse trajectory data in an objective manner, in addition to designating strongly and weakly rotating regions of a flow.

4. Conclusion and outlook

Objective coherent structure identification tools are important in a wide range of scientific and industrial fields. Methods that work equally well for gridded and sparse Lagrangian data allow meaningful comparisons between model and observations, and can provide new insights for experimental studies. In this work, we have mathematically derived new frame-independent diagnostics of trajectory rotation and stretching well-suited for coherent structure identification in sparse experimental trajectory data. Through multiple systematic comparisons with existing Lagrangian flow diagnostics, we have shown that:

- (i) \overline{rTSE} and α are objective metrics that can identify elliptic and hyperbolic LCS in highly unsteady flows, and that work equally well in flows with large or small bulk translation and rotation;
- (ii) \overline{rTSE} and α work well in many of the situations where the quasi-objective diagnostics \overline{TSE} and \overline{TRA} have previously been successful, including ocean drifter applications;
- (iii) α can designate regions with locally strong rotation in 2-D and 3-D flows for extremely sparse data;
- (iv) α values maintain accuracy and interpretability at extremely sparse sampling; and
- (v) \overline{rTSE} and α provide an avenue for meaningful LPT analysis and post-processing that does not require assumptions of steady-state flows or data interpolation to an Eulerian grid.

To assist future practitioners, we have also developed two guidelines based on the results of the present research and our experiences while conducting it.

- (i) If the trajectory concentration is below one ℓ^{-2} or ℓ^{-3} , we recommend using time-averaged values of the experiment-derived $\bar{\mathbf{v}}$, $\bar{\boldsymbol{\omega}}$ and $\bar{\mathbf{x}}$
- (ii) As the trajectory concentration decreases, the continuity of individual trajectory data becomes more important as each data point contributes more to spatial averages. It is then beneficial to select observation windows or trajectory sets for which trajectories are continuously present to prevent unphysical fluctuations in \mathbf{v}_d caused by intermittent trajectory data.

Being able to identify coherent flow features, or at least local stretching and rotation rates, using \overline{rTSE} and α for the analysis of unsteady flows in a physically meaningful, frame-indifferent manner will help advance our study of fluid dynamics in large and natural domains. Further applications of these methods, such as the time-series analytic techniques used by Aksamit *et al.* (2023), is expected to open the door for an even wider range of experiments over a vast range of scales that can harness the Lagrangian nature of turbulent structures.

Acknowledgements. The authors would like to express their gratitude to three anonymous reviewers whose comments helped improve the manuscript.

Funding. This research received no specific grant from any funding agency, commercial or not-for-profit sectors.

Declaration of interests. The authors report no conflict of interest.

Relative fluid stretching and rotation for sparse trajectories




Figure 14. Comparison of α_{ref} field for different interpolation schemes. The upper row (*a–c*) shows the α_{ref} field reconstructed using radial basis functions (rbf) for different trajectory densities. The lower row (*d–f*) shows the α_{ref} field reconstructed using standard linear interpolation for different trajectory densities. The white contours indicate the vortex boundary extracted from LAVD-field at full resolution.

Data availability statement. Matlab and Python code to perform the relative trajectory diagnostics described herein can be found on GitHub at <https://github.com/haller-group/Relative-Fluid-Stretching-and-Rotation-for-Sparse-Trajectory-Observations>. The LPT wind tunnel data that support the findings of this study are available from the authors directly. The turbulent DNS data can be found at <https://doi.org/10.15454/GLNRHK>. AVISO ocean surface current data can be found at <https://doi.org/10.48670/moi-00145>. Ocean drifter data can be accessed through <https://www.aoml.noaa.gov/phod/gdp/index.php>.

Author ORCIDs.

-  Nikolas O. Aksamit <https://orcid.org/0000-0002-2610-7258>;
-  Alex P. Encinas-Bartos <https://orcid.org/0000-0002-4203-9128>;
-  George Haller <https://orcid.org/0000-0003-1260-877X>;
-  David E. Rival <https://orcid.org/0000-0001-7561-6211>.

Appendix. Sensitivity of reconstructed α -fields to choice of interpolant

In order to reconstruct the α -field, we need to interpolate the sparse α values over a regular meshgrid. The exact values of the reconstructed α -field depend on the specific interpolation method. In this appendix, we perform a sensitivity analysis of the interpolated α -field as we progressively subsample the trajectory density in the AVISO data set (§ 3.1.2). Figure 14 compares the reconstructed α_{ref} field for different trajectory densities. For a trajectory density of $10\ell^{-2}$ the reconstructed α_{ref} field remains nearly unaffected by the interpolation scheme. For low trajectory densities ($1–0.1\ell^{-2}$), the exact boundaries of features in the α_{ref} field vary with on the choice of interpolation scheme, but we expect the overall topology of the α_{ref} field to be robust with respect to the employed interpolation method as they are the result of raw values, and are not generated by the interpolants. Specifically, local maxima of the α_{ref} clearly persist irrespective of

the chosen interpolation method. Similarly to previous studies by Encinas-Bartos *et al.* (2022), we recommend using radial basis function interpolation since it favours elliptic and smooth structures.

REFERENCES

- AKSAMIT, N.O. 2023 Mapping the shape and dimension of three-dimensional Lagrangian coherent structures and invariant manifolds. *J. Fluid Mech.* **958**, A11.
- AKSAMIT, N.O., SCHARIEN, R.K., HUTCHINGS, J.K. & LUKOVICH, J.V. 2023 A quasi-objective single buoy approach for Lagrangian coherent structures and sea ice fracture events. *Cryosphere* **17**, 1545–1566.
- BERNARD, P.S. & THOMAS, J.M. 1993 Vortex dynamics and the production of Reynolds stress. *J. Fluid Mech.* **253**, 385–419.
- BRISTOW, N., LI, J., HARTFORD, P., GUALA, M. & HONG, J. 2023 Imaging-based 3D particle tracking system for field characterization of particle dynamics in atmospheric flows. *Exp. Fluids* **64** (4), 1–14.
- BUSINGER, S., JOHNSON, R. & TALBOT, R. 2006 Scientific insights from four generations of Lagrangian smart balloons in atmospheric research. *Bull. Am. Meteorol. Soc.* **87** (11), 1539–1554.
- DEL-CASTILLO-NEGRETE, D. & MORRISON, P.J. 1993 Chaotic transport by Rossby waves in shear flow. *Phys. Fluids* **5** (4), 948–965.
- DONG, C., LIU, Y., LUMPKIN, R., LANKHORST, M., CHEN, D., MCWILLIAMS, J.C. & GUAN, Y. 2011 A scheme to identify loops from trajectories of oceanic surface drifters: an application in the Kuroshio extension region. *J. Atmos. Ocean. Technol.* **28** (9), 1167–1176.
- ENCINAS-BARTOS, A.P., AKSAMIT, N.O. & HALLER, G. 2022 Quasi-objective eddy visualization from sparse drifter data. *Chaos* **32** (11), 113143.
- FROYLAND, G. & PADBERG-GEHLE, K. 2015 A rough-and-ready cluster-based approach for extracting finite-time coherent sets from sparse and incomplete trajectory data. *Chaos* **25** (8), 087406.
- FU, S., BIWOLE, P.H. & MATHIS, C. 2015 Particle tracking velocimetry for indoor airflow field: a review. *Build. Environ.* **87**, 34–44.
- GRIFFA, A., LUMPKIN, R. & VENEZIANI, M. 2008 Cyclonic and anticyclonic motion in the upper ocean. *Geophys. Res. Lett.* **35** (1), L01698.
- GUPTA, A.K., LILLEY, D.G. & SYRED, N. 1984 *Swirl Flows*. Abacus Press.
- HADJIGHASEM, A., FARAZMAND, M., BLAZEWSKI, D., FROYLAND, G. & HALLER, G. 2017 A critical comparison of Lagrangian methods for coherent structure detection. *Chaos* **27** (5), 1–25.
- HADJIGHASEM, A., KARRASCH, D., TERAMOTO, H. & HALLER, G. 2016 Spectral-clustering approach to Lagrangian vortex detection. *Phys. Rev. E* **93** (6), 1–17.
- HALLER, G. 2015 Lagrangian coherent structures. *Annu. Rev. Fluid Mech.* **47**, 137–162.
- HALLER, G. 2023 *Transport Barriers in Flow Data: Advective, Diffusive, Stochastic and Active Methods*. Cambridge University Press.
- HALLER, G., AKSAMIT, N.O. & BARTOS, A.P.E. 2021 Quasi-objective coherent structure diagnostics from single trajectories. *Chaos* **31**, 043131.
- HALLER, G., HADJIGHASEM, A., FARAZMAND, M. & HUHN, F. 2016 Defining coherent vortices objectively from the vorticity. *J. Fluid Mech.* **795**, 136–173.
- HALLER, G. & YUAN, G. 2000 Lagrangian coherent structures and mixing in two-dimensional turbulence. *Phys. D* **147** (3–4), 352–370.
- HOU, J., KAISER, F., SCIACCHITANO, A. & RIVAL, D.E. 2021 A novel single-camera approach to large-scale, three-dimensional particle tracking based on glare-point spacing. *Exp. Fluids* **62** (5), 1–10.
- IACOBELLO, G. & RIVAL, D.E. 2023 Identifying dominant flow features from very-sparse Lagrangian data: a multiscale recurrence network-based approach. *Exp. Fluids* **64** (10), 1–14.
- JEONG, J. & HUSSAIN, F. 1995 On the identification of a vortex. *J. Fluid Mech.* **285**, 69–94.
- KAISER, F. & RIVAL, D.E. 2023 Large-scale volumetric particle tracking using a single camera: analysis of the scalability and accuracy of glare-point particle tracking. *Exp. Fluids* **64** (9), 1–12.
- KASZÁS, B., PEDERGNANA, T. & HALLER, G. 2023 The objective deformation component of a velocity field. *Eur. J. Mech. (B/Fluids)* **98**, 211–223.
- KWOK, R., CURLANDER, J.C., PANG, S.S. & MCCONNELL, R. 1990 An ice-motion tracking system at the Alaska SAR facility. *IEEE J. Ocean. Engng* **15** (1), 44–54.
- LACASCE, J.H. 2008 Statistics from Lagrangian observations. *Prog. Oceanogr.* **77** (1), 1–29.
- LEKIEN, F. & ROSS, S.D. 2010 The computation of finite-time Lyapunov exponents on unstructured meshes and for non-Euclidean manifolds. *Chaos* **20** (1), 017505.
- LEPPÄRANTA, M. 2011 *The Drift of Sea Ice*, 2nd edn. Springer.

Relative fluid stretching and rotation for sparse trajectories

- LINDSAY, R.W. & STERN, H.L. 2003 The RADARSAT geophysical processor system: quality of sea ice trajectory and deformation estimates. *J. Atmos. Ocean. Technol.* **20** (9), 1333–1347.
- LUMPKIN, R. 2016 Global characteristics of coherent vortices from surface drifter trajectories. *J. Geophys. Res.* **121** (2), 1306–1321.
- LUMPKIN, R. & CENTURIONI, L. 2019 Global Drifter Program quality-controlled 6-hour interpolated data from ocean surface drifting buoys. Dataset, <https://doi.org/10.25921/7ntx-z961>.
- LUMPKIN, R. & PAZOS, M. 2007 Measuring surface currents with Surface Velocity Program drifters: the instrument, its data, and some recent results. In *Lagrangian Analysis and Prediction of Coastal and Ocean Dynamics* (ed. A. Mariano, T. Rossby & D. Kirwan), pp. 1–23. Cambridge University Press.
- MANCHO, A.M., WIGGINS, S., CURBELO, J. & MENDOZA, C. 2013 Lagrangian descriptors: a method for revealing phase space structures of general time dependent dynamical systems. *Commun. Nonlinear Sci.* **18**, 3530–3557.
- MOWLAVI, S., SERRA, M., MAIORINO, E. & MAHADEVAN, L. 2022 Detecting Lagrangian coherent structures from sparse and noisy trajectory data. *J. Fluid Mech.* **948**, 1–34.
- ONISHCHENKO, O., FEDUN, V., HORTON, W., POKHOTILOV, O., ASTAFIEVA, N., SKIRVIN, S.J. & VERTH, G. 2021 The stationary concentrated vortex model. *Climate* **9** (3), 39.
- ROSI, G.A. & RIVAL, D.E. 2018 A Lagrangian perspective towards studying entrainment. *Exp. Fluids* **59** (1), 1–17.
- RYPINA, I.I., BROWN, M.G., BERON-VERA, F.J., KOÇAK, H., OLASCOAGA, M.J. & UDOVYDCHENKOV, I.A. 2007 On the Lagrangian dynamics of atmospheric zonal jets and the permeability of the stratospheric polar vortex. *J. Atmos. Sci.* **64** (10), 3595–3610.
- RYPINA, I.I., GETSCHER, T.R., PRATT, L.J. & MOURRE, B. 2021 Observing and quantifying ocean flow properties using drifters with drogues at different depths. *J. Phys. Oceanogr.* **51** (8), 2463–2482.
- RYPINA, I.I., SCOTT, S.E., PRATT, L.J. & BROWN, M.G. 2011 Investigating the connection between complexity of isolated trajectories and Lagrangian coherent structures. *Nonlinear Process. Geophys.* **18** (6), 977–987.
- SCHLUETER-KUCK, K.L. & DABIRI, J.O. 2017 Coherent structure colouring: identification of coherent structures from sparse data using graph theory. *J. Fluid Mech.* **811**, 468–486.
- SCHRÖDER, A. & SCHANZ, D. 2023 3D Lagrangian particle tracking in fluid mechanics. *Annu. Rev. Fluid Mech.* **55**, 511–540.
- VAN SEBILLE, E., *et al.* 2018 Lagrangian ocean analysis: fundamentals and practices. *Ocean Model.* **121** (October), 49–75.
- TAURO, F., PISCOPIA, R. & GRIMALDI, S. 2017 Streamflow observations from cameras: large-scale particle image velocimetry or particle tracking velocimetry? *Water Resour. Res.* **53** (12), 10374–10394.
- TRUESDELL, C. & NOLL, W. 2004 *The Non-Linear Field Theories of Mechanics*. 3rd edn. Springer.
- ZHOU, J., ADRIAN, R.J., BALACHANDAR, S. & KENDALL, T.M. 1999 Mechanisms for generating coherent packets of hairpin vortices in channel flow. *J. Fluid Mech.* **387**, 353–396.